

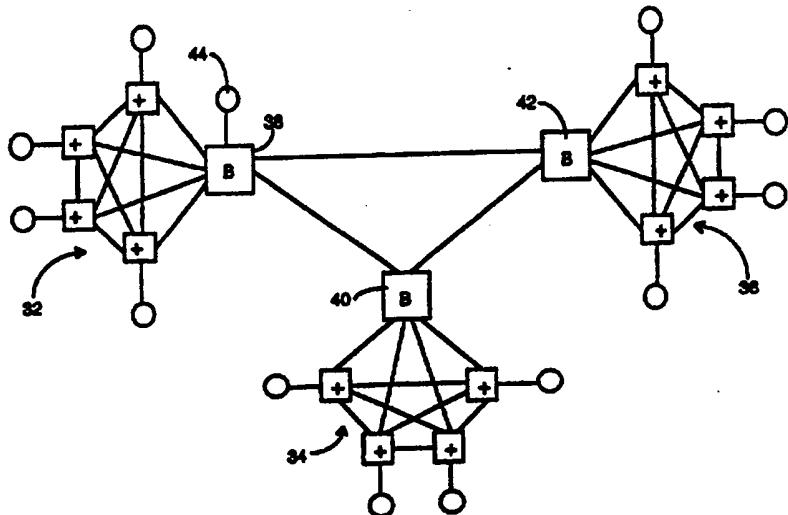
PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : H01J 13/00		A1	(11) International Publication Number: WO 96/21236
			(43) International Publication Date: 11 July 1996 (11.07.96)
(21) International Application Number: PCT/US95/16768		104 Nimbus Court, Cary, NC 27513 (US). LOWRY, David [US/US]; 2226 Camp Court, Pleasanton CA 94588 (US).	
(22) International Filing Date: 26 December 1995 (26.12.95)		(74) Agent: KREBS, Robert, E.; Burns, Doane, Swecker & Mathis, P.O. Box 1404, Alexandria, VA 22313-1404 (US).	
(30) Priority Data: 08/366,713 30 December 1994 (30.12.94) US		(81) Designated States: AL, AM, AT, AU, AZ, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TI, TM, TT, UA, UG, US, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).	
(60) Parent Application or Grant (63) Related by Continuation US 08/366,713 (CIP) Filed on 30 December 1994 (30.12.94)		Published With international search report.	
(71) Applicant (for all designated States except US): CISCO SYSTEMS [US/US]; 170 West Tasman Drive, San Jose, CA 95134-1706 (US).			
(72) Inventors; and (75) Inventors/Applicants (for US only): PERIASAMY, Ravi [IN/US]; 200 Swordgate Drive, Cary, NC 27513 (US). CLARK, Wayne [US/US]; 101 Ronsard Lane, Cary, NC 27511 (US). PANDIAN, Gnanaprakasam [IN/US]; Apartment 6, 755 North Fair Oaks, Sunnyvale, CA 94086 (US). NADERI, Ramin [US/US]; 2288 Shelley Avenue, San Jose, CA 95124 (US). BORDONARO, Frank [US/US];			

(54) Title: HIERARCHICAL ROUTER ARCHITECTURE FOR WIDE-AREA NETWORKS



(57) Abstract

A hierarchical wide-area network (30) architecture which conforms to the DLSw protocol is provided, in which multiple routers (32, 34, 36) having a logical connection to one another are designated as a peer group. At least one router in each peer group is designed as a border router (38, 40, 42), and functions as a gateway between the members of its group and other peer groups. The peer group structure reduce broadcast traffic on slow links by changing the source of the broadcast from a small access router with a relatively slow link to another, more powerful router, with a faster link on the network. More than one border router (38, 40, 42) can be included in each group, to share the transmission workload and act as a backup. The border routers can share information about the groups that they service and can cooperate to provide a single-point broadcast service to all of the users of the service. When two or more routers are connected to a network in parallel to provide back-up facilities, conditional filters are employed to prevent redundant operations. Different types of frames are handled by the two different routers so that they share the communication load between them and make the most effective use of the available bandwidth.

BEST AVAILABLE COPY

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Morocco	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

HIERARCHICAL ROUTER ARCHITECTURE FOR WIDE-AREA NETWORKS

Field of the Invention:

5 The present invention is directed to computer networks, and more particularly to an architecture for connecting several subnetworks to one another.

Background of the Invention:

When two or more computers need to communicate with one another, they are typically connected to one another by a network. Each computer that is connected to the network is referred to as a node, or station. Networks can be 10 classified according to the distances which separate the various nodes. A local-area network, or LAN, typically comprises nodes that are connected within a relatively small geographic area, such as within an office, a floor, or a building. If the nodes are separated by larger distances, for example they are in different buildings, different cities or different countries, the network which interconnects 15 them is referred to as a wide-area network, or WAN. Because of the larger distances involved, WANs usually include bridges and/or routers, which connect groups of nodes by means of telephone lines or other existing communications lines. As a result, the available bandwidth for a WAN may be significantly less than that for a LAN.

20 The different effective operating speeds of LANs and WANs can present some problems with respect to some network communications protocols. For example, one protocol that was developed for communicating over LANs is known as the Systems Network Architecture, or SNA. This protocol has certain timing constraints that must be adhered to for proper communications. For 25 example, if a source node on a network transmits an inquiry to a destination node, it expects to receive a response from the destination node within a certain period of time. If no response is received within that time, the network and/or the destination node can be interpreted to be faulty, and the network may be shut down. While the timing constraints of the SNA protocol can be easily 30 accommodated in a high speed LAN, the longer distances and the slower speeds

encountered in a WAN may delay the transmission of the response for a sufficient period of time that it is not received at the source node within the required time window. As a result, a fault condition may be indicated, when no fault may actually exist.

5 To overcome this problem, a protocol known as Data Link Switching (DLSw) was developed. Under this protocol, when a source node on one LAN sends a communication to a destination node on another LAN, a router associated with the source node reformats the information from the source node in accordance with a different protocol that is appropriate for communications over 10 WANs, such as the Transmission Control Protocol (TCP), and sends it over the network. At the receiving end, another router reformats the information in accordance with the SNA protocol, and forwards it on to the destination node. The router at the source end of the network impersonates the destination node, in response to messages sent by the source node. As a result, the source node 15 operates as if it were communicating directly with the destination node. Since this router is typically located close to the source node, the responses will be received within the proper amount of time, and therefore no fault condition is indicated.

20 In a conventional network that is configured according to the DLSw protocol, every router in the network has a direct, one-to-one connection with every other router. Depending upon the capabilities of the routers and the speeds of the transmission links over which they communicate, a network which implements the DLSw protocol is typically limited to a maximum of approximately 50 routers. As sources of information grow in number and 25 geographic disparity, this limitation of approximately 50 routers per network can impose significant constraints. For example, as a multi-national corporation expands its branch and regional offices across the world, larger networks are required to interconnect all of the offices. As a result, it is desirable to provide

a network architecture in which the potential number of routers can be significantly greater than 50.

In addition to hardware limitations, another problem that increases in magnitude as the sizes of networks grow is that of traffic congestion, particularly 5 explorer traffic. When a source node desires to send a message to a destination node, it may not always know the proper path over which to communicate with the destination node. In such a case, the source node first broadcasts an explorer frame to each of the routers connected to its network. In essence, the explorer frame inquires which routers can connect to the destination node. In response, 10 the router associated with the source node sends out appropriate inquiries to all of the other routers to which it is connected, to determine which ones have access to the destination node. In a network where each router is connected to fifty other routers, the originating router must send 50 copies of the explorer frame, each of which is addressed to a different one of the other routers. The 15 responses of these routers to the explorer frame can generate a considerable amount of traffic on the WAN. Accordingly, it is desirable to provide a network architecture which localizes broadcast frames, and thereby reduces the amount of explorer traffic in the wide-area network.

Brief Statement of the Invention:

20 In accordance with the present invention, a hierarchical wide-area network architecture which conforms to the DLSw protocol is provided, in which multiple routers having a logical connection to one another are designated as a peer group. At least one router in each peer group is designated as a border router, and functions as a gateway between the members of its group and other peer 25 groups. With this arrangement, a peer group can have up to 50 routers, and only the border router is required to contain information about, and provide a connection to, the rest of the network. To provide this function, the border router can be a more powerful one that is linked to other routers in its peer

group through a relatively fast communications link. With this arrangement, the network architecture can accommodate 2,500 or more routers.

The peer group structure can also be used to reduce broadcast traffic on slow links. This is done by changing the source of the broadcast from a small 5 access router with a relatively slow link to another, more powerful router, with a faster link on the network. One or more of these faster routers can be used to advertise group broadcast services. The smaller access routers that need to broadcast traffic to all of the routers in the group can use the services of these faster routers.

10 In an alternative embodiment of the invention, more than one border router can be included in each group, to share the transmission workload and act as a backup. Each of the two or more border routers can contain information regarding the other border routers peers, which need not be known to the other access routers that use the broadcast services. The border routers can share 15 information about the groups that they service and can cooperate to provide a single-point broadcast service to all of the users of the service. For broadcast traffic within a peer group, the two border routers can share the load among themselves.

20 In another aspect of the invention, when two or more routers are connected to a network in parallel to provide back-up facilities, conditional filters are employed to prevent redundant operations that make inefficient use of the available bandwidth. When both of the routers are operating, the conditional filters cause one of the routers to drop selected network frames, which are handled by the other router, and thereby avoid duplicate frames. Preferably, 25 different types of frames are handled by the two different routers, so that they share the communication load between them and make the most effective use of the available bandwidth.

Further features of the invention, as well as the advantages offered thereby, are explained in detail hereinafter with reference to exemplary embodiments illustrated in the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

5 Figure 1 is an architectural diagram of a conventional wide-area network which operates in accordance with the DLSw protocol;

Figure 2 is an architectural diagram of one embodiment of a wide-area network in accordance with the present invention;

10 Figure 3 is an architectural diagram of an alternate embodiment of the invention, in which two border routers are associated with a group; and

Figure 4 is an architectural diagram of a wide-area network in which two routers are connected to a network.

Detailed Description

15 To facilitate an understanding of the principles which underlie the present invention, as well as the practical applications thereof, the invention is described hereinafter in connection with certain specific embodiments. For example, reference is made to particular network protocols. It will be appreciated, however, that the applications of the invention are not limited to networks which employ only these protocols. Rather, the invention will find utility in any 20 network in which it is desirable to increase the size of the network beyond limitations imposed by hardware specifications and/or to reduce explorer traffic.

25 Referring now to Figure 1, the architecture of a typical wide-area network based upon the DLSw protocol is shown. The wide-area network comprises a number of local-area networks 10, 11, 12 and 13. Each local-area network is made up of two or more computers that are connected via a suitable configuration. For example, LAN 10 might comprise a number of computers 14, 15, 16 which are connected to one another via a token ring. Another

network 11 might consist of a number of computers 17-20 which are connected via a bus, or Ethernet connection. Each of the computers 14-20 forms a node on the network. Each of the local-area networks 10-13 has a router 22-28 associated with it. The various routers are connected by means of links which 5 form the wide-area network 30. These links might be telephone lines, for example or other dedicated lines. In essence, the wide-area network 30 connects each of the routers to every other router in the network. With this arrangement, any given node in the network can communicate with any other node.

In operation, the computers within a given local-area network 10 communicate with one another by means of a specified protocol. For example, the computers in the local-area networks 10 and 11 may each use the SNA protocol to communicate with the other computers in their given network. If a source computer 14 in one LAN 10 desires to communicate with a destination computer 17 in another LAN 11, the source computer sends a frame of 15 information to its associated router 22. This router identifies a destination router, in this case the router 24. The SNA frame of information is reformatted in accordance with a different protocol that is suitable for transmission over the wide-area network 30, for example the TCP protocol. The frame is then transmitted from the router 22 to the router 24 using this protocol. At the router 24, the received frame of information is reformatted according to the SNA 20 protocol, and then forwarded to the destination computer 17.

In an architecture of the type shown in Figure 1, there is a practical limit 25 on the number of routers that can be connected to each other via the wide-area network. For a typical router, that number is approximately fifty. Of course, it is possible to increase the total number of routers by using larger, more powerful routers. However, the cost of these more powerful routers can make the network expansion prohibitively expensive.

In accordance with the present invention, the network can be expanded to employ a larger number of routers without requiring that every router be made

more powerful. This is accomplished by means of a hierarchical peer group architecture. Referring to Figure 2, multiple routers that have a logical connection to one another, and which are expected to participate in data traffic between each other, are designated as a peer group. In the example shown in

5 Figure 2, three peer groups 32, 34 and 36 are shown. Associated with each peer group of routers is a border router 38, 40 or 42, which can be a more powerful router than the other routers of the group. The border router acts as a gateway to other peer groups. Preferably, each group consists of a number of routers that have a logical connection to one another. For example in a large corporation

10 having a number of regional offices, each of which is responsible for multiple branch offices, the peer groups can be respectively associated with the regional offices. Thus, each of the routers in the group can be respectively associated with the branch offices, and the border router for that group can be associated with the regional office responsible for those branch offices.

15 Each of the routers in a group communicates with one another. However, they do not directly communicate with the routers in any other group. Consequently, a single group can accommodate up to 50 routers itself. Group to group communications are handled via the border router within the respective groups. Thus, only the border router for a group needs to have information

20 about, and connectability to, the rest of the wide-area network. Since the border router has a greater number of connections, it is preferably a more powerful router. For example, each border router may have the ability to connect to 100 other routers. Consequently, there can be up to 50 border routers connected to one another, and each peer group can include up to 50 routers. As a result, the

25 network architecture can accommodate up to 2,500 total routers.

 In operation, all of the routers in a group are peers of one another, wherein each router can communicate directly with any other router in the group, including the border router. All of the border routers are also at a peer level with one another, i.e., there is a point to point connection from each border

router to every other border router. Whenever a router in one group desires to communicate with a router in a different group, the communications are conducted through the border routers of each of the two groups. In this type of operation, the border routers function as relay devices. In other words, the 5 border routers do not reformat, or otherwise process, the frames of data received from the other routers of their group. Rather, this data is merely passed along, or relayed, to the other border router, from where it is relayed to the destination router in its group. Of course, each border router can also have an associated local-area network. In the example of Figure 2, the border router 38 is 10 connected to a LAN 44. When one of the nodes in the LAN 44 desires to communicate with a node associated with a different group, the border router 38 can operate in the normal fashion, to reformat the data from its LAN, as required, and route it to the border router associated with the destination node. If that destination node is located on a LAN connected to the border router, that 15 router can reformat the data frame in accordance with the network protocol for its LAN, and then forward it to the destination node. Otherwise, the receiving border router functions to relay the data frame to the group router associated with the destination node.

In addition to increasing the available number of routers in the wide-area 20 network, the hierarchical group architecture of Figure 2 can also be used to reduce congestion due to broadcast traffic. For example, the individual routers of a group might typically be connected to one another by a relatively slow link, e.g., one which might have a bandwidth of 56kbps. However, in the implementation of the present invention, the more powerful border router can be 25 connected to the other routers in its peer group through a faster link, e.g., a T1 carrier having a bandwidth of 1.5 Mbps. When a router within a group needs to broadcast information to the other routers of the group, it sends the broadcast information to the border router. The border router then broadcasts the

information to all of the routers in the group, or a subset of the group, using the faster link.

In addition to providing faster broadcast services, this arrangement also reduces the number of wide-area network protocol sessions that are required on 5 the slower links. More particularly, if the wide-area network comprises a total of 2,500 routers, the originating router is not required to conduct 2,500 separate protocol sessions with each of the other routers. Rather, the originating router need only send the broadcast frame to its border router. If the border router knows the location of the destination node, it can send the frame directly to that 10 node. Otherwise, it can copy the frame and send it to all of the routers within the group, as described previously. If the destination node is associated with a different group, the border router for the originating router can send a copy of the broadcast frame to the other border routers as well, which can then forward the frame to the destination node, if known, or broadcast it to all routers within 15 the group.

In the embodiment of Figure 2, only one border router is associated with each peer group. However, it is possible to utilize two or more border routers per group. In such a case, the multiple routers can share the traffic load and function as backups for one another. Such an arrangement is illustrated in 20 Figure 3. Referring thereto, a peer group 46 has two border routers 48 and 50 associated with it. Each of these two border routers are connected with other border routers in the network. It can be seen, however, that each of the two border routers 48 and 50 are not connected to the same border routers. Rather, 25 each border router of the group is connected with a subset of the other group's border routers. In the specific example of Figure 3, the two border routers 48 and 50 of the group 46 are both connected to other border routers 52 and 54 in the network. The remaining border routers in the network, however, are only connected to one or the other of the two border routers 48 and 50.

-10-

In operation, a source router within the group 46 can choose either of the two border routers 48 or 50 to broadcast information over the network. The particular border router that is chosen can be picked arbitrarily, for example. Alternatively, the selection of the border router can be made on the basis of cost information associated with each router. For example, the source router may determine which of the two border routers provides a faster connection to a destination node, i.e., which one is connected to a faster link. If that router is available, it is chosen. If, however, much of its capacity is currently being used, it can inform the source router of this fact, in which case the other border router can be employed.

The two border routers can also cooperate with one another to provide one-point broadcast service. In particular, a source router may send a frame to one of the two border routers, e.g. the router 48, to be broadcast over the network. The border router 48 that receives this frame can send it to all of the other border routers to which it is connected. This frame is also sent to the other border router 50, which has information regarding the groups to which the first border router 48 has access. The border router 50 broadcasts the frame only to those other border routers associated with groups to which the first border 48 router does not have access.

In the embodiment of Figure 3, two routers are effectively connected to the same network, via one of the group routers. In the past, it has been known to connect two routers to the same network for backup purposes. An implementation which employs such an architecture is illustrated in Figure 4. Referring thereto, a local-area network 60 is connected to two routers 62, 64 in parallel. These two routers are connected to another router 66 via a wide-area network 68. The router 66 is connected to a second local-area network 70. In operation, if one of the routers 62 or 64 goes out of service, the other router handles the network traffic. In order for the two routers to function as backups for one another, both routers have to run in parallel, which results in a needless

waste of available bandwidth. Alternatively, the routers have to be physically powered on and off.

In accordance with another feature of the present invention, the available bandwidth provided by the two routers can be more effectively utilized, and the need to physically turn the routers on and off can be avoided, through the use of conditional filters. More particularly, each of the routers 62, 64 and 68 is provided with conditional information that determines the path to be taken when a frame is transmitted between the two LANs 60 and 70. One conditional filter is provided in the software which controls the router 68, for the condition in which a frame is transmitted from the LAN 70 to the LAN 60. Generally, the router 68 determines if one of the two parallel routers 62 and 64 is operational. For example, the router 60 may check to see whether the router 62 is operating. If so, all frames that are to be transmitted to the LAN 60 are sent to the router 62, and filtered from the router 64.

As an extension of this feature, different types of filtering can be carried out with respect to different types of information, to more effectively distribute the workload between the two routers 62 and 64. If the router 68 determines that both of the routers 62 and 64 are operational, it can send all information of one type through the router 62, and information of another type through the router 64. For example, all frames which comply with the SNA protocol can be sent to the router 62, whereas all frames conforming to a different protocol, e.g., NetBIOS, can be sent to the router 64, for transmission to the LAN 60. If the router 68 detects that either of the two routers 62 and 64 is out of service, all traffic can be sent to the other router.

A second filter is employed for information which originates in the LAN 60 and is sent to the LAN 70. In a normal operation, two copies of the same frame would be sent to the LAN 70, which are respectively transmitted through the router 62 and 64. In accordance with the present invention, however, one of the two routers can determine whether the other is operational. For example, the

-12-

router 62 can regularly send a message to the router 64, to determine whether it is up and running. If so, the router 62 can drop all frames of information from the LAN 60 to the LAN 70, so that those frames are only transmitted through the router 64. Again, the conditional filter can be employed to balance the workload 5 between the two routers. For example, the router 62 can transmit all SNA-protocol frames and drop all NetBIOS-protocol frames, whereas the router 64 can filter out all SNA frames while forwarding NetBIOS frames.

With this approach, redundant transmissions through parallel routers are avoided, and the traffic load is balanced among the various routers. However, 10 the network maintains the ability to fall back automatically to one router if the other becomes non-functional.

It will be appreciated that the conditional filtering described with reference to Figure 4 is not limited in application to the hierarchical peer group architecture of Figure 3. Rather, this filtering technique can be employed in any 15 network architecture in which two routers are physically and/or logically connected to the same network.

Preferably, the network user determines the conditions for the filters that operate in each of the routers. For example, the user can be provided with the ability to select which types of frames are to be handled by the respective 20 routers. Alternatively, the routers themselves can automatically make such a determination. For example, if one router is connected to the network by a faster link than the other router, it may be selected as the default router for all transmissions. In that case, the filter of the slower router is set to drop all frames transmitted to it by the shared network. Once the faster router reaches a 25 particular utilization level, however, it can instruct the slower router to begin to handle certain types of frames, e.g., all frames which conform to a designated protocol. The slower router then handles these frames, and the filter in the faster router is set to drop them, until such time as network utilization decreases below a threshold level.

-13-

From the foregoing, therefore, it can be seen that the present invention provides a network architecture that overcomes the size limitations previously associated with DLSw protocols. By grouping routers in a logical manner that is related to the expected communication loads among them, traffic within a group 5 can be handled more efficiently. Furthermore, the point of a broadcast message is shifted from a source node to a router, so that congestion on the network is reduced. In addition, the use of conditional filters makes more effective use of available bandwidth when two or more routers are connected to a network without losing the ability to have one router function as a backup for the other.

10 It will be appreciated by those of ordinary skill in the art that the present invention can be embodied in other specific forms without departing from the spirit or essential characteristics thereof. The presently disclosed embodiments are considered in all respects to be illustrative and not restrictive. The scope of the invention is indicated by the appended claims rather than the foregoing 15 description, and all changes that come within the meaning and range of equivalents thereof are intended to be embraced therein.

What is claimed is:

1. **A wide-area network, comprising:**
 - a plurality of groups of routers, wherein each group comprises plural routers that are logically connected to one another to provide a communication path from a network node associated with one of said routers to a network node associated with another one of said routers; and
 - 5 at least one border peer router in each group, said border peer router being logically connected to the routers in its associated groups, wherein all communications from a network node in one group to a network node in another group is transmitted through the respective border routers for said groups.
 - 10
2. **The wide-area network of claim 1 wherein the border peer router for a group can be connected to a larger number of routers than the other routers of the group.**
3. **The wide-area network of claim 1 wherein the border peer router for a group is connected to the other routers of the group by a transmission link which is faster than the transmission links which provide logical connections between said other routers of the group.**
- 15
4. **The wide-area network of claim 1 wherein at least one of said groups contains two border peer routers.**
- 20
5. **The wide-area network of claim 4 further including means in the routers of said group for selecting one of said two border peer routers for conducting communications with nodes in other groups.**
6. **The wide-area network of claim 5 wherein said selection is made on an arbitrary basis.**

-15-

7. The wide-area network of claim 5 wherein said selection is made on the basis of the cost of communicating with each of said two border peer routers.

8. The wide-area network of claim 4 further including conditional 5 filters in each of said two border peer routers for selectively transmitting or dropping frames of data that are transmitted to each of said two border peer routers from a common source.

9. The wide-area network of claim 8 further including means in each 10 of said two border peer routers for determining the operational state of the other border peer router, and for interrupting the function of the conditional filter if the other border peer router is determined to be non-operational, such that all frames of data from said common source are transmitted by the operational border peer router.

10. The wide-area network of claim 1, wherein said network conforms 15 to the Data Link Switching protocol.

11. A wide-area network, comprising:
a plurality of subnetworks, wherein each subnetwork contains at least one network node;
a plurality of routers that are respectively connected to said plurality of 20 subnetworks, for providing a logical connection from a node in one subnetwork to a node in another subnetwork, and wherein at least one of said subnetworks has at least two routers connected to it in parallel; and
a conditional filter in each of said two routers that are connected to said one subnetwork, said conditional filter causing said routers to selectively transmit or drop frames of data that are sent to said routers from said one subnetwork in 25

parallel so that only one of said two routers transmit a given frame of data from said one subnetwork.

12. The wide-area network of claim 11 wherein the conditional filter in one of said two routers causes said router to transmit frames of a first type and to drop frames of a second type, and the conditional filter in the other of said two routers causes said other router to drop frames of said first type and to transmit frames of said second type.

13. The wide-area network of claim 12 wherein said first and second types of frames respectively relate to different network transmission protocols.

10 14. The wide-area network of claim 11 further including means in each of said two routers for determining the operational state of the other router, and for interrupting the function of the conditional filter if the other router is determined to be non-operational, such that all frames of data from said one subnetwork are transmitted by the operational router.

15 15. The wide-area network of claim 11 further including means in the other routers of said wide-area network for selecting one of said two routers for conducting communications with said one subnetwork.

16. The wide-area network claim 15 wherein said selection is made on an arbitrary basis.

20 17. The wide-area network of claim 15 wherein said selection is made on the basis of the cost of communicating with each of said two routers.

-17-

18. The wide-area network of claim 11, wherein said network
conforms to the Data Link Switching protocol.

1/2

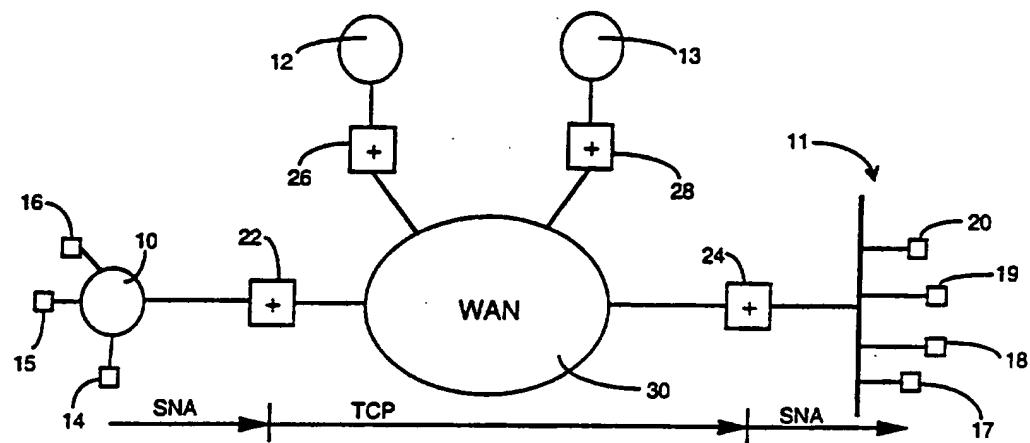


FIG. 1

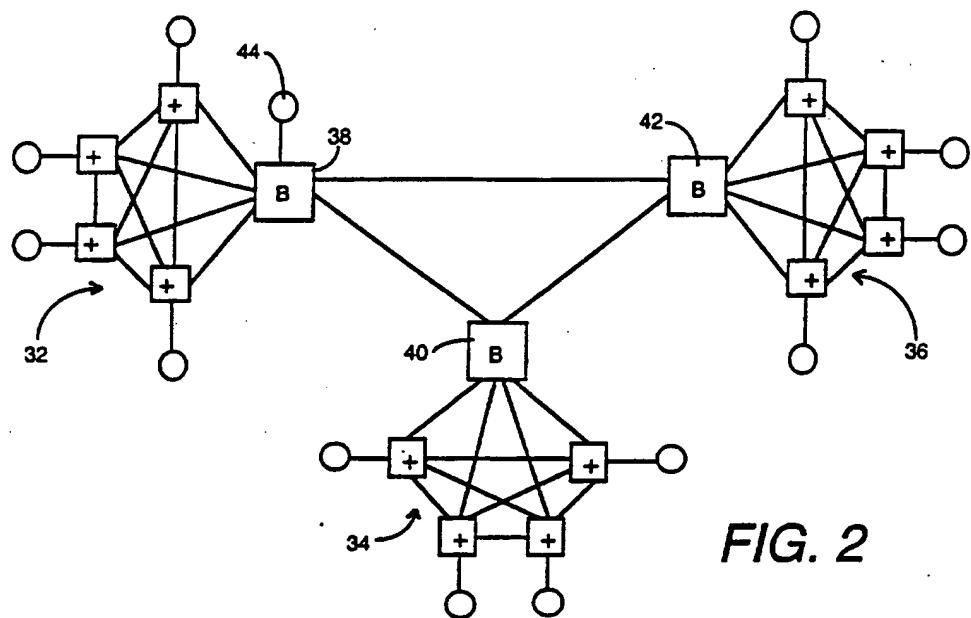
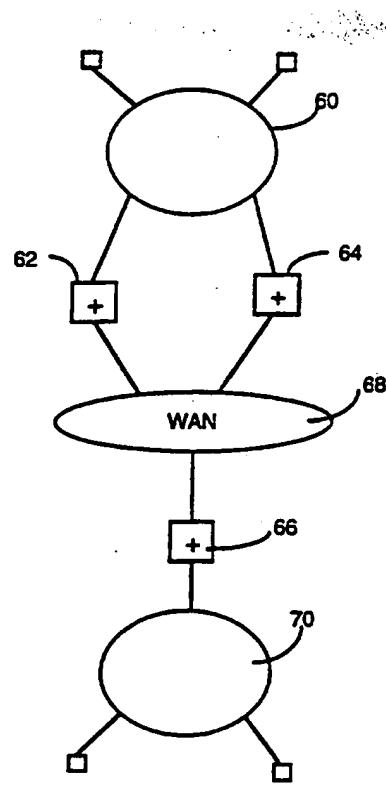
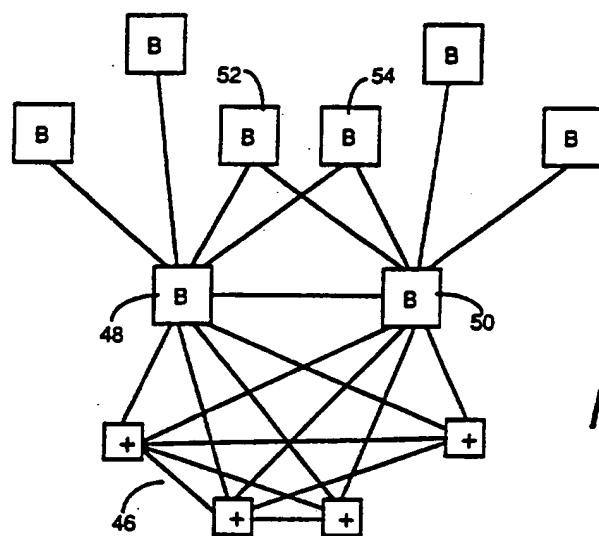


FIG. 2

2/2



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US95/16768

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : H01J 13/00

US CL :395/ 200.02,200.03,200.06, 200.15; 395/575;

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/ 200.02,200.03,200.06, 200.15; 395/575;

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS, JPOABS, DIALOG

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US, A, 5,325,362 (Aziz) 28 June 1994, abstract, cols. 1-5.	1-18
Y,P	US, A, 5,471,580 (Fujiwara et al) 28 November 1995, cols. 1-2.	1-18

 Further documents are listed in the continuation of Box C. See patent family annex.

•	Special categories of cited documents:		
"A"	document defining the general state of the art which is not considered to be part of particular relevance	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E"	earlier document published on or after the international filing date	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L"	document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O"	document referring to an oral disclosure, use, exhibition or other means	"A"	document member of the same patent family
"P"	document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

21 MARCH 1996

Date of mailing of the international search report

02 APR 1996

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer *Thomas C. Lee*
THOMAS C. LEE
Telephone No. (703) 305-9717

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.